arXiv:1511.00715v2 [cs.DS] 2 Aug 2016

The Distributed Selection Problem and the AKS Sorting Network

Piotr Berman[1] and Junichiro Fukuyama[2]

[1] Department of Computer Science and Engineering
The Pennsylvaina State University
berman@cse.psu.edu
[2] Applied Research Laboratory
The Pennsylvania State University
jxf140@psu.edu

Abstract. We consider the selection problem on a completely connected network of n processors with no shared memory. Each processor initially holds a given numeric item of b bits, allowed to send a b-bit message to another processor at a time. On such a *communication network* $\mathcal{G}$, we show that the k^{th} smallest of the n inputs can be detected in $O(\log n)$ time with $O(n \log\log n)$ messages. The possibility of such a parallel algorithm for this *distributed k-selection problem* has been unknown despite the intensive investigation on many variations of the selection problem carried out since 1970s. Satisfying the constraint of total $O(n \log\log n)$ messages, it improves on $\mathcal{G}$ the asymptotic running time of Kuhn, Locker and Wattenhofer's algorithm. Our parallel algorithm simulates the comparisons and swaps performed by the *AKS sorting network*, the n-input sorting network of logarithmic depth discovered by Ajtai, Komlós and Szemerédi in 1983. Simulation of such a network is our main trick to achieve $O(\log n)$ time and $O(n \log\log n)$ messages simultaneously. Extending its correctness proof, we will be able improve by about 47% the upper bound found by Seiferas on the constant factor of the $O(\log n)$ depth of an n-input sorting network. Furthermore, we show the universal time lower bound $\lg n$ for many basic data aggregation problems on $\mathcal{G}$. The class of problems having this lower bound is huge including the selection problem, the problem of finding the sum of n items, and that of counting items exceeding a threshold. Thus the universal lower bound means the asymptotic time optimality of our parallel algorithm.

Keywords: the selection problem, distributed selection, AKS sorting network, communication network, halver, comparator network

1 Introduction

The classical k-selection problem finds the k^{th} smallest element of given n numeric items. We consider the problem on a completely connected network $\mathcal{G}$ of n processors with no shared memory, each holding exactly one item of b bits initially. A processor node in $\mathcal{G}$ may send a b-bit message at any parallel step.

As the performance metrics, we minimize the parallel running time, and/or the total number of messages as the measure of amount of information transmitted on $\mathcal{G}$.

The parallel selection problem for connected processors with no shared memory has been extensively investigated for various network topologies, cases of how n inputs are distributed, and other constraints [1,2,3,4], as well as on the parallel comparison tree (PCT) and parallel random access machine (PRAM) models with shared memories [5,6,7]. In this paper we focus on the above case, calling it the *k-distributed selection problem* on a *communication network* $\mathcal{G}$. As suggested in [1], this case of $\mathcal{G}$ has become increasingly significant for the contemporary distributed computing applications such as sensor networks and distributed hash tables: It is common, as a metric in their performance analysis, to count the number of messages delivered at the designated destinations assuming constant time per delivery, rather than count how many times messages are forwarded by processor nodes. The considered network $\mathcal{G}$ models it well.

The distributed selection problem in this particular network case is one of the long time research topics in parallel algorithms. The results are included in the work such as [1,3,8,9]. The following summarizes only a few most closely related to our interest in the paper: Let w stand for the number of processors initially holding one or more inputs. The algorithm Frederickson and Johnson developed in [8] finds the k^{th} smallest element with $O\left(w \log \frac{k}{w}\right)$ messages on a completely connected or star-shaped processor network. Santoro et al [3] discovered in 1992 an algorithm with an expected number of messages bounded by $O\left(w \log\log \min(k, n-k) + w \log w\right)$. The more recent result in [1] explores a case when the processor network has a general diameter D; it presents a parallel algorithm with the average time bound $O\left(D \log_D n\right)$ that is shown asymptotically optimal, and one with the deterministic time bound $O\left(D \log_D^2 n\right)$.

The first contribution of this paper is a parallel algorithm for the distributed k-selection problem that runs in time $O\left(\log n\right)$ with total $O\left(n \log\log n\right)$ messages on $\mathcal{G}$. We will prove the following theorem after formulating the problem.

Theorem 1. *The k^{th} smallest of n inputs can be computed distributedly on a communication network $\mathcal{G}$ in time $O\left(\log n\right)$ with $O\left(n \log\log n\right)$ messages.* □

Such a parallel algorithm has been unknown despite the long research history on the selection problem. Satisfying the constraint of total $O\left(n \log\log n\right)$ messages, it improves the parallel time bound $O\left(D \log_D^2 n\right)$ in [1] when $\log^2 D \ll \log n$: On a processor network with diameter D, a message can be sent to anywhere forwarded by $O\left(D\right)$ processors. So the theorem means the parallel time bound $O\left(D \log n\right) \ll D \log_D^2 n$.

Our algorithm simulates the comparisons and swaps performed by the *AKS sorting network*, the n-input sorting network of $O\left(\log n\right)$ depth discovered by Ajtai, Komlós and Szemerédi in 1983 [10]. Known for the difficulty of its performance analysis, the AKS sorting network itself has been a research subject in parallel algorithm since then. Paterson [11] simplified its construction. Seiferas [12] further clarified it with the estimate that the depth can be at most

$7 \cdot 6.07 \lg n = 48.79 \lg n$ layers of 1/402.15-halvers. Here an *ε-halver* $\left(\varepsilon \in \left(0, \frac{1}{2}\right)\right)$ is a comparator network defined in [11,13], such that for every positive integer $z \leq \frac{m}{2}$, the left half of the output includes at most εz items among the z largest inputs, and the right half at most εz among the z smallest inputs. In the recent work by Goodrich [14], the constant factor of the total $O(n \log n)$ nodes is significantly reduced although its depth bound is $O(n \log n)$. The overall research interest on the AKS sorting network has been simpler understanding and reduction of the constant factor of the depth.

In order to design our parallel algorithm for the distributed selection problem, we show that if the depth of the AKS sorting network is reduced from $\Theta(\log n)$ to $\Theta(\log \log n)$, the number of items in the wrong half of the output is $O\left(n \log^{-c} n\right)$ for any given constant $c > 0$. The parallel algorithm simulates it, being our main trick to achieve $O(\log n)$ time and $O(n \log \log n)$ messages simultaneously. Formally we will show:

Theorem 2. *For each $c \in \mathbb{R}^+$ and sufficiently large power n of 2, there exists a comparator network $\mathcal{H}$ to re-order n inputs satisfying the following three.*

i) *$\mathcal{H}$ is an AKS sorting network of reduced depth.*
ii) *$\mathcal{H}$ is a weak $\lg^{-c} n$-halver, i.e., the first half of the ordered n outputs includes at most $n \lg^{-c} n$ items larger than the median, and the second half at most $n \lg^{-c} n$ smaller than the median.*
iii) *There are no more than $22.34c \lg \lg n$ layers of 1/175-halvers in $\mathcal{H}$.* □

Here an *AKS sorting network of reduced depth* means that the lower nodes of depth more than $O(\log \log n)$ are simply removed. We will define it exactly when we construct $\mathcal{H}$ in Section 3. The proof of Theorem 2 is the second contribution of the paper.

Thirdly, we will show a smaller constant factor of the $O(\log n)$ depth of the n-input sorting network. Extending the proof of Theorem 2, we will confirm:

Theorem 3. *For each sufficiently large $n \in \mathbb{Z}^+$, there exists a sorting network on n inputs with at most $25.54 \lg n$ layers of 1/395-halvers.* □

This improves the above Seiferas's estimate by about 47%.

Our fourth result is the running time lower bound $\lg n$ for many data aggregation problems on $\mathcal{G}$. The class of problems having the lower bound is huge including the selection problem and many others. The following statement will be confirmed as a corollary to the theorem we will show in Section 6.

Corollary 1. *Any parallel algorithm takes at least $\lg n$ steps in the worst case to compute each of the following three problems on a communication network distributedly: i) the k-selection problem on n inputs each of at least $\lceil \lg n \rceil + 1$ bits; ii) the problem of finding the sum of n inputs; iii) the problem of counting items among n inputs, each exceeding a given threshold.* □

By our formulation in Section 2, the statement assumes that each of the n processor nodes holds an input at time 0. The corollary shows the asymptotic time optimality of our parallel algorithm as well.

The rest of the paper consists as follows. In Section 2, we define general terminology showing related facts. We will prove Theorem 2 in Section 3, which is extended to our improved constant estimate in Section 4. Theorem 1 is verified in Section 4. We prove the universal lower bound $\lg n$ in Section 5, followed by concluding remarks in Section 6.

2 General Terminology and Related Facts

In this paper, a *communication network* $\mathcal{G}$ means a collection of n processor nodes each two of which are connected, capable of exchanging a b-bit message at a parallel step. We consider a computational problem P whose input is a set of n numeric items of b bits *distributed over* $\mathcal{G}$, *i.e.*, each processor holds exactly one item at time 0. A parallel algorithm A is said to *compute* P *on* $\mathcal{G}$ *distributedly in time* t *with* m *messages*, if all the bits of the computed result from n inputs distributed over $\mathcal{G}$ is stored at a designated processor node after the t^{th} parallel step with total m messages. Messages may be exchanged asyncronously on $\mathcal{G}$.

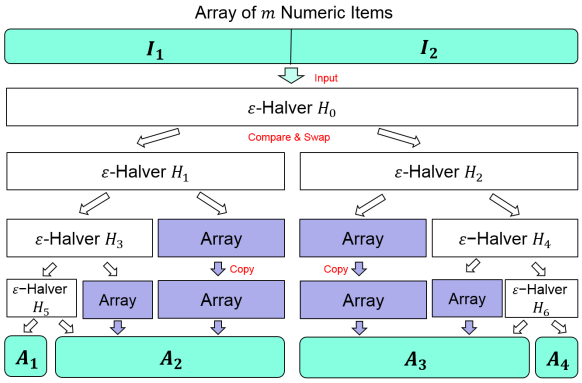

Fig. 1. A Separator S in the AKS Sorting Network

A *comparator network* $\mathcal{H}$ is a directed acyclic graph consisting of *comparators* as nodes of positive depth, and nodes of depth 0 each storing an input numeric item. Such a comparator of depth $d \geq 1$ can receive two items from nodes of depth $d-1$ sending their maximum and/or minimum to other comparators of depth $d+1$. It can thus *swap* two items, or *copy* by receiving a same item. For notational convenience, we say that the *width* of $\mathcal{H}$ is the maximum number of nodes with a same depth. In the standard terminology, a numeric item input to a comparator is called *wire*. We may call an ordered set of wires *array*, for which the set theoretical notation is used.

The AKS sorting network is a comparator network to sort n inputs, whose depth is $\Theta(\log n)$ and width no more than n, performing $O(n \log n)$ comparisons.

We observe here that the k^{th} smallest of n items distributed over $\mathcal{G}$ can be detected in $O(\log n)$ time with $O(n \log n)$ messages the following way. We design such a parallel algorithm A so that a processor node v in $\mathcal{G}$ simulates from time ct to $c(t+1)$ (c : a constant, $t \in \mathbb{N}$) a comparator of depth t in the AKS sorting network $\mathcal{H}$. We mean by "simulate" that v receives two items from other nodes to send their minimum and/or maximum to anywhere in $\mathcal{G}$. Due to the width of the AKS sorting network, and since $\mathcal{G}$ is completely connected, n processors in $\mathcal{G}$ can simulate all the swaps and copying performed by $\mathcal{H}$. This way A sorts the n inputs in $O(\log n)$ steps, then fetches the k^{th} smallest one. We can therefore achieve $O(\log n)$ time and $O(n \log n)$ messages to compute the distributed k-selection problem on $\mathcal{G}$.

Our parallel algorithm improves the above so that the number of messages is reduced to $O(n \log \log n)$. To modify the AKS sorting network, we look into the details of its unit component called *separator* given in a textbook [13]: A separator illustrated as S in Fig. 1 consists of the seven ε-halvers $H_0, H_1, \ldots, H_6$, for a sufficiently small constant $\varepsilon \in (0, 1)$. Let I_1 and I_2 be the left and right halves of the input array to S, respectively. The separator S re-orders $I_1 \cup I_2$ into four disjoint output arrays A_1, A_2, A_3 and A_4 such that

$$|A_1| = |A_4| = \frac{1}{16}m, \quad \text{and} \quad |A_2| = |A_3| = \frac{7}{16}m.$$

The way numeric items are compared in each ε-halver is represented by a (d_0, ε)-*expander graph* $\left(d_0 \in \mathbb{Z}^+\right)$. It is a d_0-regular bipartite graph (V_1, V_2, E) such that $|V_1| = |V_2|$ and

$$\varepsilon|\Gamma(U)| \geq (1-\varepsilon)\min(\varepsilon|V_i|, |U|), \tag{1}$$

for every subset U of V_1 or V_2, where $\Gamma(U)$ is the neighbor set of U. The input to the halver is an array of $|V_1 \cup V_2|$ elements where each comparison is represented by an edge in E. When two items are compared, they are swapped to correct the order if necessary. We can choose ε arbitrarily small by increasing d_0 to a sufficiently large constant.

The seven components $H_0, H_1, \cdots, H_6$ of S all satisfy (1), meeting the condition to be an ε-halver. Suppose the left output of H_0 contains a set U of more than εz elements each at least the z^{th} largest element of $I_1 \cup I_2$. By (1), $|\Gamma(U)| \geq (1-\varepsilon)z > z - |U|$, meaning that $\Gamma(U)$ includes an element less than the z^{th} largest input. A contradiction against correct swapping. So the left output of H_0 includes at most εz elements among the z largest inputs.

We also have the facts below. We will use them in our proof in Section 3.

Lemma 1. *Let S be a separator on m inputs with the seven ε-halvers as in Fig. 1, where $m \in 16\mathbb{Z}^+$, and $\varepsilon \in \left(0, \frac{1}{100}\right)$ is a constant. The following three hold true for every $g \in \mathbb{R}$:*

i) Suppose there are l items larger than g in I_1, and r items not exceeding g in I_2. Then $A_1 \cup A_2$ includes at most $\frac{1}{2}\varepsilon m + \max(l - r, 0)$ items larger than g, and $A_3 \cup A_4$ includes at most $\frac{1}{2}\varepsilon m + \max(r - l, 0)$ items not exceeding g.
ii) Suppose the A_4 output includes p items larger than g where $p \leq \frac{m}{32}$. Then $A_1 \cup A_2 \cup A_3$ includes at most $\frac{4\varepsilon p}{1-4\varepsilon}$ items larger than g.
iii) If the A_1 output includes $p' \leq \frac{m}{32}$ items not exceeding g, then $A_2 \cup A_3 \cup A_4$ includes at most $\frac{4\varepsilon p'}{1-4\varepsilon}$ items not exceeding g.

Proof. i): WLOG assume $r \geq l$. The top halver H_0 of S swaps so that there are $r - y$ items larger than g in $A_1 \cup A_2$, and $l - y$ items not exceeding g in $A_3 \cup A_4$, for some $y \geq 0$. It suffices to verify $l - y \leq \frac{1}{2}\varepsilon m$ to have i). Suppose not. Then there are more than $\frac{1}{2}\varepsilon m$ elements of $A_1 \cup A_2$ larger than g, and more than $\frac{1}{2}\varepsilon m$ of $A_3 \cup A_4$ not exceeding g.

Let U be the set of those elements of $A_1 \cup A_2$. By (1), $|\Gamma(U)| \geq \frac{1-\varepsilon}{\varepsilon} \min\left(\frac{\varepsilon m}{2}, |U|\right) > \frac{(1-\varepsilon)m}{2}$. So the items in U would be compared more than $\frac{m}{2} - \frac{1}{2}\varepsilon m$ elements of $A_3 \cup A_4$, which contradicts the last sentence of the first paragraph. The contradictions prove i).

ii), iii): We only show ii) as the proof of iii) is similar. Let z be the number of items larger than g in $I_1 \cup I_2$. We claim that $z \leq \left(\frac{1}{32} + 2\varepsilon\right) m$. By the same argument as above, the right output of the halver H_0 includes at least $\min\left(\frac{m}{2}, z\right) - \frac{1}{2}\varepsilon m$ items larger than g. Also those of H_2, H_4 and H_6 include at least $\min\left(\frac{m}{4}, z\right) - \varepsilon m$, $\min\left(\frac{m}{8}, z\right) - \frac{3}{2}\varepsilon m$ and $\min\left(\frac{m}{16}, z\right) - 2\varepsilon m$ such elements, respectively. So $\min\left(\frac{m}{16}, z\right) - 2\varepsilon m \leq p \leq \frac{m}{32}$, meaning $\frac{m}{32} \geq z - 2\varepsilon m$. The claim follows this.

Because of $z \leq \left(\frac{1}{32} + 2\varepsilon\right) m$, and since H_0 is an ε-halver, its left output includes at most εz items larger than g. Considering H_2, H_4 and H_6 similarly, $A_1 \cup A_2 \cup A_3$ contains at most $4\varepsilon z$ such items. Its ratio to p is no more than $\frac{4\varepsilon}{1-4\varepsilon}$, proving ii). □

In Section 4, we will have a situation where the input size $m = |I_1 \cup I_2|$ may not be a multiple of 16. We modify the ε-halvers in S in such a case.

Construction of a separator S for any input size $m \in \mathbb{Z}^+$: If m is odd, add an extra element, which is compared with no other elements in H_0, to the left input of H_0, and remove it from the output. Similarly, add an extra element to the right input to H_1, H_3 or H_5 if the input size is odd. Add an extra element to the left input to H_2, H_4 or H_6 if the input size is odd. Remove them from their output arrays.

As a result,

$$\frac{m}{16} - \frac{1}{16} \leq |A_1| \leq \frac{m}{16} + \frac{7}{8}, \quad \frac{7m}{16} - \frac{21}{16} \leq |A_2| \leq \frac{7m}{16},$$
$$\frac{m}{16} \leq |A_4| \leq \frac{m}{16} + \frac{15}{16}, \quad \text{and} \quad \frac{7m}{16} - \frac{7}{8} \leq |A_3| \leq \frac{7m}{16} + \frac{7}{16}. \tag{2}$$

The upper bound in the first case is due to $|A_1| \leq \frac{1}{2}\left(\frac{1}{2}\left(\frac{1}{2}\left(\frac{m}{2}+1\right)+1\right)+1\right)$. The other bounds are found similarly.

For this S, the three statements of Lemma 1 can be shown true if we replace ε by

$$\hat{\varepsilon} = \varepsilon + \frac{300}{299}\varepsilon^2, \tag{3}$$

in i) when $\varepsilon = 1/395$. We see it as a corollary to the lemma.

Corollary 2. *Let $m \in \mathbb{Z}^+$, $\varepsilon = 1/395$, and S be a separator on m inputs with the seven ε-halvers. The three statements i), ii) and iii) of Lemma 1 hold true if ε is replaced by $\hat{\varepsilon}$ in i).*

Proof. i): By the same argument as the lemma, we can show the bounds $\frac{m+1}{2} + \max(l-r, 0)$ and $\frac{m+1}{2} + \max(r-l, 0)$ for the left and right halves, respectively. If $m \geq \frac{299}{300\varepsilon}$, then $\varepsilon(m+1) \leq \varepsilon\left(m + \frac{300}{299}\varepsilon m\right) = \hat{\varepsilon}m$, so i) holds in this case. If $m < \frac{299}{300\varepsilon}$, then $\varepsilon(m+1) \leq \frac{299}{300} + \varepsilon < 1$. The number $\varepsilon(m+1)$ in place of εm is an upper bound for an integral value, so the bounded value is zero. Thus i) holds in every case with $\hat{\varepsilon}$.

ii), iii): We saw $\min\left(\frac{m}{16}, z\right) - 2\varepsilon m \leq p \leq \frac{m}{32}$ in the lemma. Instead we have

$$\min\left(\frac{m}{16} - \frac{1}{16}, z\right) - 2\varepsilon(m+1) \leq p.$$

Here $-2\varepsilon m$ is replaced by $-2\varepsilon(m+1)$ since the input size of any H_i does not exceed $m+1$. Also $\frac{m}{16}$ is replaced by $\min(|A_1|, |A_4|) \geq \frac{m}{16} - \frac{1}{16}$.

Thus $\min\left(\frac{m}{16} - \frac{1}{16}, z\right) - 2\varepsilon(m+1) \leq p \leq \frac{m}{32}$, meaning $z \leq \left(\frac{1}{32} + 2\hat{\varepsilon}\right)m$: If $\frac{m}{16} - \frac{1}{16} - 2\varepsilon(m+1) \leq \frac{m}{32}$, then $m \leq 2$ so $p \leq \frac{m}{32}$ is zero. Otherwise $z \leq \frac{m}{32} + 2\varepsilon(m+1) \leq \left(\frac{1}{32} + 2\hat{\varepsilon}\right)m$, seen similarly to i).

The remaing argument is the same as the lemma. The array $A_1 \cup A_2 \cup A_3$ includes at most $4\varepsilon z$ items larger than g. Its ratio to p does not exceed $\frac{4\varepsilon}{1-4\varepsilon}$. The statement ii) follows. Similarly for iii). □

3 The AKS Sorting Network as a Weak Halver

We show Theorem 2 in this section. Below we construct such an AKS sorting network $\mathcal{H}$ of reduced depth $O(\log\log n)$. Our main task is to bound the number of items in the wrong half of the output array, called *strangers*, by $n\lg^{-c} n$ where the constant $c > 0$ is given by the theorem. Lemma 3 below will lead to the depth bound, *i.e.*, at most $22.34c\lg\lg n$ layers of 1/175-halvers.

Assume without loss of generality that n inputs have distinct values[3]. There are $j_{max}+1$ *layers* in $\mathcal{H}$, where

$$j_{max} = \frac{c}{\lg(4/e)}\lg\lg n,$$

[3] For s elements of a same value g, regard that they are valued $g+\delta, g+2\delta, \ldots, g+s\delta$ for an infinitesimal number $\delta > 0$. All the statements in our proof remain true with this change.

with $e = 2.71828\ldots$ being the natural logarithm base. In this section, obvious ceiling/floor functions are omitted. Each layer is a complete binary tree of separators of depth $d_{max} - j$ where

$$d_{max} = \frac{c}{5} \lg\lg n + j_{max} = c\left(\frac{1}{\lg(4/e)} + \frac{1}{5}\right) \lg\lg n.$$

Every separator consists of the seven ε-halvers as in Fig. 1 where

$$\varepsilon = \frac{1}{175}.$$

Also we create two arrays L and R that will be the left and right halves of the output array of $\mathcal{H}$, respectively.

Denote by $S_{d,j}$ $(0 \le j \le j_{max},\ 0 \le d \le d_{max} - j)$ a separator node of depth d in the j^{th} layer of $\mathcal{H}$. Here are the main construction rules.

The Construction Rules of the AKS Sorting Network $\mathcal{H}$ of Reduced Depth:

- If a separator $S_{d,j}$ $(d \ge 1,\ j \le j_{max} - 1)$ is the left child of its parent $S_{d-1,j}$, the A_1 and A_4 outputs of $S_{d,j}$ are sent to the left of the copy of $S_{d-1,j}$ in the $j + 1^{st}$ layer. That is, the sent array is concatenated with the first element of the destination array. Also if $S_{d,j}$ is the right child of $S_{d-1,j}$, its $A_1 \cup A_4$ is sent to the right of the copy of $S_{d-1,j}$ in the the $j + 1^{st}$ layer.
- The A_2 and A_3 outputs of $S_{d,j}$ such that $d < d_{max} - j$ are sent to its left and right children in the same layer, resp.
- The A_2 and A_3 outputs of $S_{d_{max}-j,j}$ for every j are sent to L and R, resp.
- The A_1 and A_4 outputs of $S_{0,j}$ such that $j < j_{max}$ are sent to the left and right of $S_{0,j+1}$, resp.
- The A_1 and A_4 outputs of $S_{d,j_{max}}$ for all d are sent to L and R, resp.
- The union $L \cup R$ of the finally obtained L and R is the output of $\mathcal{H}$ as a weak $\lg^{-c} n$-halver.

As mentioned in Section 1, we call the comparator network $\mathcal{H}$ constructed by this set of rules an *AKS sorting network of reduced depth.* It is identical with the standard AKS sorting network given in [13] except for its depth and the two output arrays L and R.

We introduce extra terminology. A *stranger* is a numeric item greater than the median of n inputs to $\mathcal{H}$ existing in the left half of $\mathcal{H}$, or one less than or equal to the median existing in the right half. Define

$m_{d,j}$: the number of inputs to each $S_{d,j}$, *i.e.*, $|I_1 \cup I_2|$ of $S_{d,j}$.

$p_{d,j}$: the number of strangers existing in $\bigcup_{i=1}^{4} A_i$ of all $S_{d,j}$, and

$$\Delta = \frac{1}{32}.$$

We observe the following four.

a) The number of layers of ε-halvers in $\mathcal{H}$ does not exceed

$$4d_{max} + 8j_{max} = c\left(\frac{12}{\lg(4/e)} + \frac{4}{5}\right)\lg\lg n < 22.34c\lg\lg n,$$

as desired. It is confirmed by the first lemma below.

b) We have

$$\left(\frac{7}{16}\right)^d m_{0,j} \le m_{d,j} \le n\left(\frac{7}{16}\right)^d 8^{-j}\frac{(d+2j)^j}{j!},$$

for each $j \le j_{max}$ and $d \le d_{max}$, regarding $0^0 = 0! = 1$. The lower bound is due to $m_{d,j} \ge \frac{7}{16}m_{d-1,j}$, seen by construction. The upper bound is proved in Lemma 3 below.

c) By Stirling's approximation [15], $j! > \left(\frac{j}{e}\right)^j$, so b) means

$$m_{d,j} \le \left(\frac{7}{16}\right)^d \left(\frac{e}{8}\left(2 + \frac{d}{j}\right)\right)^j n,$$

when $j \ge 1$. Thus $m_{0,j} \le \left(\frac{e}{4}\right)^j n$.

d) Our invariant for $j \le j_{max}$ and $d \le d_{max} - j$ is

$$p_{d,j} \le \Delta^{d+1} m_{0,j}.$$

We prove it in the third lemma below. With this and c), the number of strangers in $L \cup R$ is upper-bounded by

$$\sum_{j=0}^{j_{max}} p_{d_{max}-j,j} + \sum_{d=0}^{d_{max}-j_{max}} p_{d,j_{max}}$$
$$< 2\Delta^{d_{max}-j_{max}+1} n + 2\Delta\left(\frac{e}{4}\right)^{j_{max}} n < n\lg^{-c} n,$$

as $j_{max} = \frac{c}{\lg(4/e)}\lg\lg n$ and $d_{max} - j_{max} = \frac{c}{5}\lg\lg n = \frac{c}{-\lg\Delta}\lg\lg n$.

By a) and d), the comparator network $\mathcal{H}$ satisfies all the desired properties. It remains to prove the three lemmas mentioned above. We will have Theorem 2 after it.

Lemma 2. *Let $h_{d,j}$ be the number of layers of ε-halvers in the sub-network of $\mathcal{H}$ rooted at $S_{d,j}$. Then $h_{d,j} \le 4\,(d_{max} - d + 2j_{max} - 2j)$ for every $j \le j_{max}$ and $d \le d_{max} - j$.*

Proof. Proof by double induction on j and d. We show the general induction step first. Assume true for $j+1$ and all d, and for j and $d+1$. Prove true for j and d. By induction hypothesis, $h_{d-1,j+1} \le 4\,(d_{max} - (d-1) + 2j_{max} - 2(j+1))$ and $h_{d+1,j} \le 4\,(d_{max} - (d+1) + 2j_{max} - 2j)$. By construction, the depth $h_{d,j}$ does not exceed their maximum plus 4, as the separator $S_{d,j}$ has four layers of ε-halvers. Thus, $h_{d,j} \le 4\,(_{max} - d + 2j_{max} - 2j)$, proving the induction step for $j \ge 1$ and $d \ge 1$.

The claim is clearly true when $j = j_{max}$, and can be shown similarly for the case $j < j_{max}$ and $d = d_{max}$. The lemma follows. □

Lemma 3.

$$m_{d,j} \leq n \left(\frac{7}{16}\right)^d 8^{-j} \frac{(d+2j)^j}{j!},$$

for every $j \leq j_{max}$ *and* $d \leq d_{max} - j$.

Proof. Proof by double induction on j and d. The basis $j = d = 0$ is trivially true as $0^0 = 0! = 1$. Assume true for $j = 0$ and $d - 1$. The desired bound holds for $j = 0$ and d due to $m_{d,0} = \frac{7}{16} m_{d-1,0}$ and the induction hypothesis. Thus it is true for $j = 0$ and all d.

Assume true for $j-1$ and prove true for j and $d = 0$. We have the recurrence

$$m_{0,j} = \frac{1}{8} m_{0,j-1} + \frac{1}{4} m_{1,j-1},$$

by construction. Find with the induction hypothesis that

$$j(2j-2)^{j-1} + \frac{7j}{8}(2j-1)^{j-1} \leq (2j)^j$$

suffices. The inequality holds true since when $j \geq 2$, we have $(2j)^j \geq (2j-1)^j + j(2j-1)^{j-1}$ and $(2j-1)^j \geq (2j-2)^j + j(2j-2)^{j-1}$ by the binomial theorem. This proves the lemma for $j \geq 1$ and $d = 0$.

Assume true for $j-1$ and all d, and for j and $d-1$. Prove true for j and d. We similarly find from $m_{d,j} = \frac{1}{8} m_{d+1,j-1} + \frac{7}{16} m_{d-1,j}$ that $\frac{7}{16} j(d+2j-1)^{j-1} + (d+2j-1)^j \leq (d+2j)^j$ suffices. It is true by the binomial theorem again. This completes the induction step. □

Lemma 4. $p_{d,j} \leq \Delta^{d+1} m_{0,j}$ *for each* $j \leq j_{max}$ *and* $d \leq d_{max} - j$.

Proof. Proof by double induction on j and d. For the basis $j = d = 0$, we observe that the number of strangers in I_1 of $S_{0,0}$ equals that in I_2. By Lemma 1 i), the separator $S_{0,0}$ reduces strangers to at most $\varepsilon m_{0,0} < \Delta m_{0,0}$, proving the basis.

Assume true for $j = 0$ and $d - 1$ and prove true for $j = 0$ and d. Consider each $S_{d,0}$. There are at most $\Delta^d m_{0,0}$ strangers in $A_1 \cup A_4$ of its parent $S_{d-1,0}$ by induction hypothesis. This number does not exceed $\Delta m_{d-1,0} \geq \Delta \left(\frac{7}{16}\right)^{d-1} m_{0,0}$. By Lemma 1 ii) or iii), if p strangers are output from $A_1 \cup A_4$ of $S_{d,j-1}$, its $A_2 \cup A_3$ includes at most $\frac{4\varepsilon}{1-4\varepsilon} p$ strangers.

Consider all $S_{d,0}$ and $S_{d-1,0}$. By induction hypothesis and the above, there are no more than $\frac{4\varepsilon}{1-4\varepsilon} p_{d-1,0} \leq \frac{4\varepsilon}{1-4\varepsilon} \Delta^d m_{0,0} \leq \Delta^{d+1} m_{0,0}$ strangers in $A_2 \cup A_3$ of all $S_{d-1,0}$. (This holds for $\varepsilon = 1/175$ and $\Delta = 1/32$.) So $p_{d,0} < \Delta^{d+1} m_{0,0}$, proving the induction step. We have the lemma for $j = 0$ and all d.

Assume true for $j-1$ and prove true for j and $d = 0$. By induction hypothesis, $p_{d',j-1} \leq \Delta^{d'+1} m_{0,j-1}$ for all $d' \leq d_{max} - j + 1$. So there are at most

$$\sum_{d'=2}^{d_{max}-j+1} p_{d',j-1} < 2\Delta^3 m_{0,j-1}$$

strangers *not* input to $S_{0,j}$. This bounds the difference between the numbers of strangers in I_1 and I_2 of $S_{0,j}$. By Lemma 1 i), $S_{0,j}$ reduces strangers to at most

$$\varepsilon m_{0,j} + 2\Delta^3 m_{0,j-1} \le \varepsilon m_{0,j} + 16\Delta^3 m_{0,j} \le \Delta m_{0,j},$$

where $m_{0,j} \ge \frac{1}{8} m_{0,j-1}$ by construction. This proves the lemma for j and $d = 0$.

Assume true for $j - 1$, and j and $d - 1$. Prove true for j and d. Consider each $S_{d,j}$ and its parent $S_{d-1,j}$. Suppose $S_{d-1,j}$ outputs p strangers from $A_1 \cup A_4$. By induction hypothesis, $p \le \Delta^d m_{0,j}$ that does not exceed $\Delta m_{d-1,j} \ge \Delta \left(\frac{7}{16}\right)^{d-1} m_{0,j}$. Lemma 1 again assures that $A_2 \cup A_3$ of $S_{d-1,j}$ includes at most $\frac{4\varepsilon}{1-4\varepsilon} p$ strangers. Therefore,

$$p_{d,j} \le p_{d+1,j-1} + \frac{4\varepsilon p_{d-1,j}}{1-4\varepsilon} \le \Delta^{d+2} m_{0,j-1} + \frac{4\varepsilon}{1-4\varepsilon} \Delta^d m_{0,j} \le \Delta^{d+1} m_{0,j},$$

by induction hypothesis and $m_{0,j} \ge \frac{1}{8} m_{0,j-1}$. This verifies the induction step. □

The lemmas complete the proof of Theorem 2.

4 Improved Constant Factor of the $O(\log n)$ Depth of an n-Input Sorting Network

Modify the AKS sorting network $\mathcal{H}$ constructed in Section 3 so that

$$j_{max} = \left\lceil \frac{\lg n}{\lg(4/e)} \right\rceil, \quad d_{max} = \lg n + j_{max} + 6, \quad \text{and} \quad \varepsilon = \frac{1}{395}, \tag{4}$$

while $\Delta = 1/32$ remains the same. Create separators $S_{d,j}$ for all j and d such that $0 \le j \le j_{max}$ and $0 \le d \le d_{max}$ rather than $d \le d_{max} - j$. With this setting for sufficiently large n, the depth of $\mathcal{H}$ as a graph of ε-halvers does not exceed

$$4\left(d_{max} + 2j_{max} + 1\right) \le 4\left(1 + 3\lg^{-1}\frac{4}{e}\right)\lg n + 18 < 25.54 \lg n,$$

seen similarly to Lemma 8. In this section, we prove that the AKS sorting network with the setting sorts the n inputs correctly. This will confirm Theorem 3 improving the aforementioned estimate in [12] that the depth can be at most $7 \cdot 6.07 \lg n = 48.79 \lg n$ layers of 1/402.15-halvers.

Due to the new depth of $\mathcal{H}$, the input array sizes $m_{d,j}$ of some separators $S_{d,j}$ may not be divisible by 16. Construct $S_{d,j}$ by the rule given below Lemma 1. The changes affect $\mathcal{H}$ as follows.

- The input sizes of two $S_{d,j}$ may differ due to varying rounding errors generated by construction. For simplicity, we regard $m_{d,j}$ as $|I_1 \cup I_2|$ of the currently considered separator $S_{d,j}$.

- Lemma 3 is modified into $m_{d,j} \leq n \left(\frac{7}{16}\right)^d 8^{-j} \frac{(d+2j)^j}{j!} + 13$, proved as Lemma 9 in Appendix. When $j \geq 1$,

$$m_{d,j} \leq \left(\frac{7}{16}\right)^d \left(\frac{e}{8}\left(2+\frac{d}{j}\right)\right)^j n + 13, \tag{5}$$

  by Stirling's approximation.
- It means $m_{d,j_{max}} \leq 14$ for any d: In $m_{d,j} \leq \left(\left(\frac{7}{16}\right)^{\frac{d}{j}} \frac{e}{8}\left(2+\frac{d}{j}\right)\right)^j n + 13$, we have $\left(\frac{7}{16}\right)^{\frac{d}{j}} \frac{e}{8}\left(2+\frac{d}{j}\right)$ approaching zero as $\frac{d}{j}$ grows larger, and not exceeding $\frac{e}{4}$. So $m_{d,j_{max}} \leq \left(\frac{e}{4}\right)^{j_{max}} n + 13 \leq 14$ since $j_{max} \geq \frac{\lg n}{\lg(4/e)}$.

With these we see:

Lemma 5. *All the n items input to $\mathcal{H}$ are also input to each layer.*

Proof. Observe the following.
- By construction, $|A_4| \geq 1$ if $m_{d,j} \geq 1$, and $|A_1| \geq 1$ if $m_{d,j} \geq 2$ for every $S_{d,j}$.
- $m_{d,j} \leq 13$ if $d \geq \lg n + j$. It is true for $j = 0$ by Lemma 9. When $j \geq 1$, by (5), $m_{d,j} - 13$ is at most

$$\left(\frac{7}{16}\right)^d \left(\frac{e}{8}\left(2+\frac{d}{j}\right)\right)^j n \leq n 2^{-\lg n} \left(\frac{e}{2\cdot 8}\left(2+\frac{d}{j}\right)\left(\frac{7}{8}\right)^{\frac{d}{j}}\right)^j < 1,$$

  whose second inequality is due to $\frac{e}{16}(2+x)\left(\frac{7}{8}\right)^x < 1$ for every $x \geq 0$.
- So $|A_1 \cup A_4| \leq 2$ if $d \geq \lg n + j$.
- The above two mean $m_{d,j} \leq 10$ for every $S_{d,j}$ such that $d \geq \lg n + j + 1$: The separator receives at most $\left\lceil \frac{13-2}{2} \right\rceil$ elements from the parent and 4 elements from the two $S_{d+1,j-1}$.
- Likewise, $m_{d,j} \leq 8$ if $d \geq \lg n + j + 2$, and $m_{d,j} \leq 7$ if $d \geq \lg n + j + 3$.

We show by induction on j that:
i) $m_{d,j} \leq 3$ if $d = \lg n + j + 4$,
ii) $m_{d,j} \leq 1$ if $d = \lg n + j + 5$, and
iii) $m_{d,j} = 0$ if $d \geq \lg n + j + 6$.

We only show the induction step as the basis $j = 0$ is proved similarly. Consider $S_{d,j}$ such that $d \geq \lg n + j + 4$. We have $m_{d-1,j} \leq 7$ from the above, and $m_{d+1,j-1} \leq 0$ by induction hypothesis. So $m_{d,j} \leq \left\lceil \frac{7-2}{2} \right\rceil = 3$, proving i). The other two are shown similarly.

The lemma follows iii) since $d_{max} = \lg n + j_{max} + 6$. □

Define arrays associated with a considered separator $S_{d,j}$ the following way. Fix a given input array to $\mathcal{H}$. Denote by S_L and S_R the left and right children of $S_{d,j}$, respectively. Let $B(S_{d,j})$ be the array recursively defined by

$$B(S_{d,j}) = (A_1 \text{ of } S_{d,j}) \cup B(S_L) \cup B(S_R) \cup (A_4 \text{ of } S_{d,j}),$$

if $d < d_{max}$. It equals $\bigcup_{i=1}^{4} A_i$ of $S_{d,j}$ if $d = d_{max}$.

It can be seen by induction on j that $B(S_{d,j})$ consists of the items input to the subtree rooted at $S_{d,j}$ in the j^{th} layer. With Lemma 5, we define $B\left(S_{0,j_{max}}\right)$ as the output of $\mathcal{H}$.

Consider the complete binary tree of separators in the j^{th} layer. Initially, let I be the sorted array of the n inputs. We say that it is *correct for* $S_{0,j}$. Recursively, a sorted array I correct for a separator $S_{d,j}$ is given so that $|I| = |B\left(S_{d,j}\right)|$. Split I into the four sub-arrays of sizes $|A_1|$, $|B\left(S_L\right)|$, $|B\left(S_R\right)|$, and $|A_4|$ in the order. They are correct for A_1, S_L, S_R and A_4, respectively.

We generalize the definition of strangers. A numeric item x is said to be an *l-stranger* if there exists $S_{l,j}$ satisfying one of the following two:

I) x is in the array correct for A_1 (A_4) of $S_{l,j}$ or S_L (S_R), but is incorrectly output from A_4 (A_1) of $S_{l,d}$ or its descendant.
II) x is in the array correct for A_1 or A_4 of $S_{l,j}$ but is incorrectly output from a proper descendant of $S_{l,d}$.

The following statement is our invariant.

Lemma 6. *For each separator $S_{d,j}$ and index $l \leq d$, there are $\Delta^{d-l+1} m_{d,j}$ or less l-strangers output from $S_{d,j}$.* □

Its proof is found in Appendix. The lemma means that $B\left(S_{0,j_{max}}\right)$ is sorted, proving that $\mathcal{H}$ satisfying (4) correctly re-orders the n inputs. Verify it as follows: Let $j = j_{max}$. There are no strangers in the j^{th} layer by Lemma 6 and $m_{d,j} \leq 14$. This implies that $B\left(S_{d,j}\right)$ for every $S_{d,j}$ consists of the elements in the array correct for $S_{d,j}$. It also means that $B(S_{d,j})$ is sorted if both $B\left(S_L\right)$ and $B\left(S_R\right)$ are sorted. For a leaf $S_{d_{max},j}$, the array $B\left(S_{d_{max},j}\right)$ is empty thus sorted, by iii) in the proof of Lemma 5. Hence $B\left(S_{0,j}\right)$ is sorted, completing our proof of Theorem 3.

5 Selection in $O\left(\log n\right)$ Time with $O\left(n \log \log n\right)$ Messages on a Communication Network

We prove Theorem 1 in this section. Our parallel algorithm for the distributed k-selection problem on a communication network $\mathcal{G}$ uses the weak halver $\mathcal{H}$ given by Theorem 2 and three other AKS sorting networks as sub-components. Keep assuming that n is a large power of 2 and the input array has elements of distinct values. Let

$$q = \lg n.$$

Choose $c = 2$ for the constant c in Theorem 2, so $\mathcal{H}$ is a weak $n \lg^{-2} n$-halver. All the expressions of array sizes in this section omit the floor or ceiling function. We first describe the algorithm focusing on the case $k = n/2$.

Algorithm 1 for Finding the Median of n Inputs:

1. Apply $\mathcal{H}$ to the given n inputs. Let L and R be its left and right halves of the output, respectively, and $C = \emptyset$.
2. Arbitrarily partition L into $\frac{n}{2q}$ subarrays each of q items naming them $L_1, L_2, \ldots, L_{\frac{n}{2q}}$.
3. Find the maximum value l_i in each L_i.
4. Sort the $\frac{n}{2q}$ items $l_1, l_2, \ldots, l_{\frac{n}{2q}}$ by an AKS sorting network. Add the elements of L_i to C for every i such that l_i is among the largest $n \lg^{-2} n$ of $l_1, l_2, \ldots, l_{\frac{n}{2q}}$.
5. Perform Steps 2–4 to R symmetrically.
6. Sort the elements of C by another AKS sorting network. Find its median, returning it as the median of all n.

Each AKS sorting network used above re-orders at most $2n \lg^{-1} n$ inputs. This allows us to detect the median with a sufficiently small number of comparisons, leading to $O(n \log \log n)$ messages. We confirm the correctness of the algorithm first.

Lemma 7. *Algorithm 1 correctly finds the median of all n inputs.*

Proof. We claim that no element of $L - C$ is a stranger or the median of n after Step 4. Let i be an index such that l_i is not among the largest $n \lg^{-2} n$ of $l_1, l_2, \ldots, l_{\frac{n}{2q}}$. If an element of L_i were a stranger or the median, the largest $n \lg^{-2} n$ items would be all strangers, contradicting Theorem 2. This proves the claim.

By symmetry, no element of $R - C$ is a stranger. Steps 4 and 5 remove the same number of items from the both halves, each neither a stranger nor the median. The array C includes the same number of strangers in the both halves. Hence Step 6 correctly finds the median of all. □

We now describe how to run Algorithm 1 on the communication network $\mathcal{G}$. Find our implementation below noting two remarks.

- Regard that there are $2n$ processor nodes in $\mathcal{G}$ instead of n, since any of them at odd and even time slots can play two different roles.
- Each processor in $\mathcal{G}$ can simulate any of the four comparator networks the way mentioned in Section 2: A processor receives two numeric items from other nodes to send their minimum and/or maximum to anywhere in $\mathcal{G}$.

Implementation of Algorithm 1 on $\mathcal{G}$: Simulate two AKS sorting networks in Steps 4 and 5 with extra n processors. Then select the elements of C in $O(\log \log n)$ time as follows: Initially the first node of each L_i is notified if l_i is among the largest $n \lg^{-2} n$. The jth inductive step ($j = 1, 2, \ldots, \lg q - 1$) merges two adjacent sets of items in C. Before the jth step, two sets of $2^{j-1}|L_i|$ nodes hold each at most $2^{j-1}|L_i|$ items in $L \cap C$, in such a way that the smallest numbered nodes hold one item in each. During the jth step the 2nd set of nodes move items to the 1st by sending messages if necessary. After it, the elements in $L \cap C$ are held by the first consecutive nodes of the $2^j|L_i|$ nodes.

In the final step, $L \cap C$ is merged with $R \cap C$ the same way. Perform Step 6 with a simulated AKS sorting network on the obtained C. This completes the description of our implementation.

Algorithm 1 correctly runs on $\mathcal{G}$ in $O(\log n)$ time with $O(\log \log n)$ messages. One can check it with Lemma 7 noting that:

- The simulated $\mathcal{H}$ runs in $O(\log \log n)$ time with $O(n \log \log n)$ messages.
- Each of the three simulated AKS sorting network runs in $O(\log n)$ time with $O(n)$ messages.
- In the process to construct C, the jth step sends $O(n)$ messages in total to move items between adjacent sets of nodes.

This verifies Theorem 1 for $k = \frac{n}{2}$.

Selecting an item but the median is done similarly. Assume without loss of generality that we want the kth smallest item such that $k < \frac{n}{2}$. We add extra $n - 2k$ elements valued $-\infty$ to the input array. Then the median of the new array is the kth smallest of the original. Algorithm 1 detects it with the same asymptotic bounds on the running time and number of messages.

We have constructed a parallel algorithm on $\mathcal{G}$ that distributedly computes the kth smallest of the n inputs in $O(\log n)$ time with $O(n \log \log n)$ messages. We now have Theorem 1.

6 The Universal Parallel Time Lower Bound $\lg n$ on a Communication Network

In this section, we show that it takes at least $\lg n$ steps to compute many basic data aggregation problems on a communication network $\mathcal{G}$ including the distributed selection problem. Consider a parallel algorithm to compute a function $f(x_1, x_2, \ldots, x_n) \in \{0, 1\}$ where $x_1, x_2, \ldots, x_n$ are input numeric items of b-bits distributed over $\mathcal{G}$. Such f is said to be *critical everywhere* if there exist $x_1, x_2, \ldots x_n$ such that

$$f(x_1, x_2, \ldots, x_{i-1}, \hat{x}_i, x_{i+1}, \ldots, x_n) \neq f(x_1, x_2, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n), \quad (6)$$

for each index $i = 1, 2, \ldots, n$ and some b-bit numeric item $\hat{x}_i$ depending on i.

Below we prove the time lower bound $\lg n$ to compute such f as the following theorem.

Theorem 4. *Let f be a function on n inputs critical everywhere. It takes at least $\lg n$ parallel steps in the worst case to compute f on a communication network distributedly.* □

By the theorem, we will have the statement mentioned in Section 1.

Corollary 1. Any parallel algorithm takes at least $\lg n$ steps in the worst case to compute each of the following three problems on a communication network distributedly: i) the k-selection problem on n inputs each of at least $\lceil \lg n \rceil + 1$ bits; ii) the problem of finding the sum of n inputs; iii) the problem of counting items among n inputs, each exceeding a given threshold.

Proof. i): Given such an algorithm on n inputs $x_1, x_2, \ldots, x_n$, let $f(x_1, x_2, \ldots, x_n)$ be the least significant bit of the k^{th} smallest input. Such a function f is critical everywhere: Choose the n inputs such that $x_i = 2(i-1)$. If $k > 1$, let $\hat{x}_i = x_k - 1$ for every $i \leq n$, otherwise $\hat{x}_i = 1$. (All of these numbers are $\lceil \lg n \rceil + 1$-bit integers.) These satisfy (6). By the theorem, the algorithm takes at least $\lg n$ steps to compute such f on a communication network distributedly.

ii): Shown similarly to i) by choosing f as the least significant bit of the sum of n inputs.

iii): Similarly by choosing f as the least significant bit of the number of inputs exceeding a given threshold. □

Consider the class of problems to compute on a communication network distributedly such that any particular bit of the computed result is a function critical everywhere. It is very large containing many aggregation problems as the above three, each with the parallel running time lower bound $\lg n$ on $\mathcal{G}$.

We prove the theorem. Denote by v_i the i^{th} processor node of the communication network $\mathcal{G}$, and by x_i the numeric item held by v_i at time 0. Also let A be a parallel algorithm to compute f on $\mathcal{G}$ distributedly. Since f is critical everywhere, there exists an input set $X = \{x_1, x_2, \ldots, x_n\}$ such that (6). Assume $f(x_1, x_2, \ldots, x_n) = 1$ without loss of generality, written as $f(X) = 1$.

We focus on the computation performed by A on the input set X. For time $t \geq 0$ and index $i = 1, 2, \ldots, n$, define the set $V_{i,t}$ of nodes in $\mathcal{G}$ recursively as follows.

- $V_{i,0} = \{v_i\}$ for time $t = 0$.
- $V_{i,t} = V_{i,t-1} \cup V_{j,t-1}$ if v_j sends a message to v_i at time $t \geq 1$. Otherwise $V_{i,t} = V_{i,t-1}$.

$V_{i,t}$ is the set of processors v_j such that the values of x_j can possibly affect the computational result at v_i at time t.

Let t_{max} be the time when A terminates on X, and v_i be the processor that decides $f(X) = 1$ at time t_{max}. The size of $V_{i,t_{max}}$ must be n; otherwise A decides $f(X) = 1$ at v_i with no information on some input x_j, contradicting that f is critical everywhere. The following lemma formally confirms it.

Lemma 8. $|V_{i,t_{max}}| = n$.

Proof. By the equivalence between an algorithm and Boolean circuit [16], there exists a Boolean circuit C to compute $f(X)$, which is converted from A by the reduction algorithm. Construct the subgraph of C that decides $f(X) = 1$ at v_i as follows: For each $t = 0, 1, \ldots, t_{max}$, with the algorithm running on X at each processor v_j and the exchanged messages, inductively construct a Boolean circuit to decide each bit of the computed result stored at v_j at time t. For the time $t = t_{max}$ and processor v_i, we have a circuit C' that decides $f(X) = 1$ only from x_j such that $v_j \in V_{i,t_{max}}$.

The existence of C' means $f(X) = 1$ if all x_j such that $v_j \in V_{i,t_{max}}$ have the values specfied by X. If $|V_{i,t_{max}}| < n$, this contradicts that f is critical everywhere satisfying (6). Hence $|V_{i,t_{max}}| = n$. □

We also have $|V_{j,t}| \leq 2^t$ for each time t and index j. It is because $|V_{j,t}| = |V_{j,t-1} \cup V_{l,t-1}| \leq 2 \cdot 2^{t-1} = 2^t$ if a processor v_l sends a message to v_j at time $t \geq 1$.

Threfore, $\lg n = \lg |V_{i,t_{max}}| \leq t_{max}$. Theorem 4 follows.

7 Concluding Remarks and Open Problems

We have shown $25.54 \lg n$ layers of $1/395$-halvers as an upper bound on the depth of the AKS sorting network, and $22.34c \lg \lg n$ layers of $1/175$-halvers as that on the depth of a weak $\lg^{-c} n$-halver. The obtained bounds can be further improved by choosing another parameter set. It is a question if we can significantly reduce these values. Regarding the number $O(n \log n \log n)$ of messages to find the k^{th} smallest item in $O(\log n)$ time, it remains open whether or not it is asymptotically optimal on a communication network. We conjecture it positively.

Appendix: Proof of (5) and Lemma 6

We have (5) when $j \geq 1$ as a result of Lemma 9 below. As in Section 4, $m_{d,j}$ denotes the input size $|I_1 \cup I_2|$ of the currently considered separator $S_{d,j}$ of depth d in the j^{th} layer of the AKS sorting network $\mathcal{H}$ satisfying (4).

Lemma 9. *$m_{d,j} \leq n \left(\frac{7}{16}\right)^d 8^{-j} \frac{(d+2j)^j}{j!} + 13$ for each $j \leq j_{max}$ and $d \leq d_{max}$.*

Proof. Denote by $m^*_{d,j}$ the maximum of $m_{d,j}$ for all $S_{d,j}$. It satisfies the following recursive relation.

$$m^*_{d,j} \leq \begin{cases} \frac{m^*_{d+1,j-1}}{4} + \frac{7m^*_{d-1,j}}{16} + \frac{65}{16}, & \text{if } j \geq 1 \text{ and } d \geq 1. \\ \frac{m^*_{0,j-1}}{8} + \frac{m^*_{1,j-1}}{4} + \frac{87}{16}, & \text{if } j \geq 1 \text{ and } d = 0. \\ \frac{7m^*_{d-1,j}}{16} + \frac{7}{16}, & \text{if } j = 0 \text{ and } d \geq 1. \\ n, & \text{if } j = d = 0. \end{cases}$$

We have the upper bound in the first case since

$$m^*_{d,j} \leq 2\left(\frac{m^*_{d+1,j-1}}{8} + \frac{29}{16}\right) + \left(\frac{7m^*_{d-1,j}}{16} + \frac{7}{16}\right) \leq \frac{m^*_{d+1,j-1}}{4} + \frac{7m^*_{d-1,j}}{16} + \frac{65}{16}.$$

It is true by (2), since it means $|A_1|+|A_4|$ of $S_{d+1,j-1}$ does not exceed $\frac{m_{d+1,j-1}}{8} + \frac{29}{16}$. Find the other inequalities similarly.

To verify the lemma, we show by double induction on j and d that

$$m^*_{d,j} \leq n\left(\frac{7}{16}\right)^d 8^{-j} \frac{(d+2j)^j}{j!} + 13. \tag{7}$$

We first prove the general induction step. Assume true for $j-1$ and all d, and j and $d-1$. Prove true for j and d.

The inequality in the first case above is equivalent to

$$m^*_{d,j} - 13 \le \frac{1}{4}\left(m^*_{d+1,j-1} - 13\right) + \frac{7}{16}\left(m^*_{d-1,j} - 13\right).$$

By induction hypothesis, $m^*_{d+1,j-1} - 13 \le \left(\frac{7}{16}\right)^{d+1} 8^{-j+1} \frac{(d+2j-1)^{j-1}}{(j-1)!}$, and $m^*_{d-1,j} - 13 \le \left(\frac{7}{16}\right)^{d-1} 8^{-j} \frac{(d+2j-1)^j}{j!}$. Similarly to the proof of Lemma 3, we find that it suffices to verify $\frac{7}{16} j(d+2j-1)^{j-1} + (d+2j-1)^j \le (d+2j)^j$, which is true by the binomial theorem. This proves (7) for $j \ge 1$ and $d \ge 1$.

We show (7) for $j \ge 1$ and $d = 0$ assuming true for $j-1$. With $m^*_{0,j} - 13 \le \frac{1}{8}\left(m^*_{0,j-1} - 13\right) + \frac{1}{4}\left(m^*_{1,j-1} - 13\right)$ derived from the second case of the recursive relation, we find that $j(2j-2)^{j-1} + \frac{7j}{8}(2j-1)^{j-1} \ge (2j)^j$ suffices. It is true since $(2j)^j \ge (2j-1)^j + j(2j-1)^{j-1}$ and $(2j-1)^j \ge (2j-2)^j + j(2j-2)^{j-1}$, proving (7) for $j \ge 1$ and $d = 0$.

(7) is true for $j = d = 0$ by $0^0 = 0! = 1$, and shown similarly for $j = 0$ and $d \ge 1$. This completes the proof. □

The following lemma is a preparation to verify Lemma 6.

Lemma 10. *Let $m_{d,j}$ and $m_{d-1,j}$ be the input sizes of a separator $S_{d,j}$ and its parent $S_{d-1,j}$, respectively. Then $m_{d,j} \le m_{d-1,j} + 14$.*

Proof. Proof by induction on $j \le j_{max}$. Assuming true for $j-1$, we first show the induction step by inner induction on d.

We show the basis $d = 1$ of inner induction. For the given $S_{d,j}$, consider its child $S_{1,j}$ and the copy $S_{1,j-1}$ in the $j-1^{st}$ layer. From (2),

$$\begin{aligned} m_{0,j} &\ge \left(\frac{1}{8} m_{0,j-1} - \frac{1}{16}\right) + 2\left(\frac{1}{8}\left(\frac{7}{16} m_{0,j-1} - \frac{21}{16}\right) - \frac{1}{16}\right), \quad \text{and} \\ m_{1,j} &\le \left(\frac{7}{16} m_{0,j} + \frac{7}{16}\right) + 2\left(\frac{1}{8}\left(\frac{7}{16} m_{1,j-1} + \frac{7}{16}\right) + \frac{29}{16}\right). \end{aligned} \tag{8}$$

Also

$$m_{0,j} \le \left(\frac{1}{8} m_{0,j-1} + \frac{29}{16}\right) + 2\left(\frac{1}{8}\left(\frac{7}{16} m_{0,j-1} + \frac{7}{16}\right) + \frac{29}{16}\right) = \frac{15}{64} m_{0,j-1} + \frac{355}{64}.$$

Substitute this into the second line of (8), from which we find that

$$m_{1,j} - m_{0,j} - 14 \le \frac{7}{64}\left(m_{1,j-1} - m_{0,j-1} - 14\right).$$

By induction hypothesis, the both hand sides are negative, proving the basis $d = 1$.

Assume true for $d-1$ and prove true for d. For the given $S_{d,j}$ and $S_{d-1,j}$, let $S_{d-2,j}$, $S_{d-1,j-1}$ and $S_{d,j-1}$ be the parent of $S_{d-1,j}$, the copy of $S_{d-1,j}$ in the

$j-1^{st}$ layer, and that of $S_{d,j}$, resp. For these separators,

$$m_{d-1,j} \geq 2\left(\frac{1}{8}\left(\frac{7}{16}m_{d-1,j-1} - \frac{21}{16}\right) - \frac{1}{16}\right) + \left(\frac{7m_{d-2,j}}{16} - \frac{21}{16}\right), \quad \text{and}$$
$$m_{d,j} \leq 2\left(\frac{1}{8}\left(\frac{7}{16}m_{d,j-1} + \frac{7}{16}\right) + \frac{29}{16}\right) + \frac{7m_{d-1,j}}{16} + \frac{7}{16},$$

from which we find

$$m_{d,j} - m_{d-1,j} - 14 < \frac{7}{64}\left(m_{d,j-1} - m_{d-1,j-1} - 14\right) + \frac{7}{16}\left(m_{d-1,j} - m_{d-2,j} - 14\right) \leq 0,$$

proving the induction step.

It remains to show the claim when $j = 0$. It is similarly proved by inner induction on d with $m_{d-1,0} \leq \frac{7}{16}m_{d,0} + \frac{7}{16}$. The lemma follows. □

We now prove our invariant with

$$\Delta = \frac{1}{32}, \quad \varepsilon = \frac{1}{395}, \quad \text{and} \quad \hat{\varepsilon} = \varepsilon + \frac{300}{299}\varepsilon^2 < \frac{1}{393},$$

from (3).

Lemma 6. For each separator $S_{d,j}$ and index $l \leq d$, there are $\Delta^{d-l+1}m_{d,j}$ or less l-strangers output from $S_{d,j}$, *i.e.*, in $\bigcup_{i=1}^{4} A_i$ of $S_{d,j}$.

Proof. We show the lemma by induction on j, l and d. We first show the general induction step for $j \geq 1$ and $l \geq 1$. Assume true for $j-1$ and all l, and for j and $1, 2, \ldots, l-1$. Prove true for j and l by inner induction on d.

The basis of the inner induction occurs when $d = l$. Consider each $S_{d,j}$. Let α be the total number of l'-strangers ($0 \leq l' \leq l$) sent to all the proper descendants $S_{d',j}$ of $S_{d,j}$ from the $j-1^{st}$ layer. By induction hypothesis,

$$\alpha \leq \sum_{\substack{d+1 \leq d' \leq d_{max} \\ 0 \leq l' \leq l = d}} 2^{(d'+1)-d}\Delta^{(d'+1)-l'+1}m^*_{d'+1,j-1}, \tag{9}$$

where $m^*_{d'+1,j-1}$ denotes the maximum value of $m_{d+1,j-1}$ under consideration (*i.e.*, of all $S_{d'+1,j-1}$ sending items to $S_{d',j}$).

Defining $m^*_{d',j}$ similarly, we have $m^*_{d',j} \geq \frac{1}{8}m^*_{d'+1,j-1} - \frac{1}{16}$ from (2), and $m^*_{d',j} \leq m_{d,j} + 14(d'-d)$ by Lemma 10. So

$$m^*_{d'+1,j-1} \leq 8\left(m_{d,j} + 14(d'-d) + \frac{1}{16}\right) \leq 141m_{d,j} \cdot 1.35^{d'-d}: \tag{10}$$

If $m_{d,j} \geq 1$, the right inequality is implied by $141m \cdot 1.35^x \geq 8(m + 14x + 1/16)$ for every $m \geq 1$ and $x \geq 0$. Otherwise $m_{d,j} = 0$ so $S_{d,j}$ outputs no l-strangers.

By (9) and (10),

$$\begin{aligned}\alpha &\leq \sum_{\substack{d+1 \leq d' \leq d_{max} \\ 0 \leq l' \leq l = d}} 2.7^{d'-d}\Delta^{d'-l'} \cdot 282\Delta^2 m_{d,j} \\ &< \frac{2.7\Delta}{1-2.7\Delta} \cdot \frac{282\Delta^2}{1-\Delta}m_{d,j} < 0.839\Delta m_{d,j},\end{aligned} \tag{11}$$

holding true for $\Delta = 1/32$.

Let β be the number of l''-strangers ($0 \le l'' \le l-1$) input to $S_{l,j}$. We claim that

$$\beta \le \left(1 + \frac{4\varepsilon}{1-4\varepsilon}\right) \frac{\Delta^2}{1-\Delta} m_{d,j} < 0.0326 \Delta m_{d,j}. \tag{12}$$

By induction hypothesis, there are $\sum_{l''=0}^{l-1} \Delta^{d-l''+1} m_{d,j} < \frac{\Delta^2}{1-\Delta} m_{d,j}$ or less l''-strangers output from $S_{d,j}$. Denote by I the array correct for $S_{d,j}$, by p_1 the number of l''-strangers in $A_1 \cup A_4$ exceeding the maximum in I, and by p_2 the number of those less than the minimum in I. By ii) and iii) of Corollary 2 with $p_i < \frac{\Delta^2}{1-\Delta} m_{l,j} < \Delta m_{l,j}$, the total number of l''-strangers input to $S_{d,j}$ is no more than $(p_1+p_2)\left(1 + \frac{\varepsilon}{1-4\varepsilon}\right) < \left(1 + \frac{4\varepsilon}{1-4\varepsilon}\right) \frac{\Delta^2}{1-\Delta} m_{d,j}$, proving (12). Note that this counts all the l''-strangers satisfying I) or II) in Section 4.

To complete the basis $d = l$, we will apply Corollary 2 i) to $I_1 \cup I_2$ of $S_{d,j}$. Set g as the $|I_1|^{th}$ smallest element of I. Let q_1 and q_2 be the number of elements in I_1 greater than g and that of elements in I_2 not exceeding g, resp. Then

$$|q_1 - q_2| \le \alpha + \beta : \tag{13}$$

All that may contribute to $|q_1 - q_2|$ are l''-strangers input to the subtree rooted at $S_{d,j}$ (*i.e.*, the elements not belonging to I as $l'' < l = d$), and l-strangers sent to the proper descendants of $S_{d,j}$ from the $j-1^{st}$ layer. Their total number is bounded by $\alpha + \beta$.

Hence, by (11), (12), (13) and Corollary 2 i), there are

$$|q_1 - q_2| + \hat{\varepsilon} m_{l,j} \le \alpha + \beta + \hat{\varepsilon} m_{l,j} \le \Delta m_{l,d}$$

or less l-strangers output from $S_{d,j}$. This proves the basis $d = l$ of the inner induction.

We show the induction step. Assume true for $d-1$ and prove true for d. Consider each $S_{d,j}$ and its parent $S_{d-1,j}$. There are $\Delta^{(d-1)-l+1} m_{d-1,j} \le \Delta m_{d-1,j}$ or less l-strangers in $A_1 \cup A_4$ of $S_{d-1,j}$ by induction hypothesis. By Corollary 2 ii) and iii), and also (2), the number of l-strangers sent from $S_{d-1,j}$ to $S_{d,j}$ is upper-bounded by

$$\frac{4\varepsilon}{1-4\varepsilon} \Delta^{d-l} m_{d-1,j} \le \frac{4\varepsilon}{1-4\varepsilon} \Delta^{d-l} \cdot \frac{16}{7} \left(m_{d,j} + \frac{21}{16}\right) \le \frac{9.152\varepsilon}{1-4\varepsilon} \Delta^{d-l} m_{d,j} :$$

If $\frac{4\varepsilon}{1-4\varepsilon} \Delta m_{d-1,j} < 1$, its left hand size is zero. So $m_{d-1,j} \ge \frac{1-4\varepsilon}{4\Delta\varepsilon} \ge 3128$, meaning $m_{d,j} \ge \frac{7}{16} m_{d-1,j} - \frac{21}{16} > 1367$. The second inequality above is derived from $\frac{16}{7}\left(m_{d,j} + \frac{21}{16}\right) < \frac{16}{7}\left(1 + \frac{1}{1367} \cdot \frac{21}{16}\right) m_{d,j}$.

Also there are no more than

$$\Delta^{(d+1)-l+1} \left(m^1_{d+1,j-1} + m^2_{d+1,j-1}\right) \le 8\Delta^{d-l+2} \left(m_{d,j} + \frac{1}{8}\right) \le 8.008 \Delta^{d-l+2} m_{d,j}$$

l-strangers sent from the $j-1^{st}$ layer. Here $m^i_{d+1,j-1}$ $(i = 1, 2)$ denotes the input sizes of the two $S_{d+1,j-1}$ sending items to $S_{d,j}$. Assume $\Delta^2 \left(m^1_{d+1,j-1} + m^2_{d+1,j-1}\right) \geq 1$ to derive the second inequality similarly[4].

As a result, the number of l-strangers output from $S_{d,j}$ is upper-bounded by

$$\frac{9.152\varepsilon}{1-4\varepsilon}\Delta^{d-l}m_{d,j} + 8.008\Delta^{d-l+2}m_{d,j} \leq \Delta^{d-l+1}m_{d,j},$$

true with $\varepsilon = 1/395$ and $\Delta = 1/32$. This completes the induction step for d, as well as that for j and l.

It remains to show the claim for the base cases. The arguments are all similar to the general induction step. One can prove it with minor changes noting that:

- For $j \geq 1$ and $l = 0$, the differences from the above are that $S_{0,j}$ receives items from $S_{0,j-1}$ instead of $S_{l-1,j}$, and that there are no l'-strangers such that $l' < l$.
- When $j = l = d = 0$, the claim holds by Corollary 2 i), since there are the same number of 0-strangers in both I_1 and I_2 of $S_{0,0}$. This requires $\hat{\varepsilon} \leq \Delta$.
- For the other cases of $j = 0$, the difference from the above is that there are no nodes in the $j - 1^{st}$ layer sending items to $S_{d,0}$. This requires $\frac{4\varepsilon}{1-4\varepsilon} \leq \Delta$.

The lemma follows. □

References

1. Kuhn, F., Locher, T., Wattenhofer, R.: Tight bounds for distributed selection. In: Proc. of the 19th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), pp. 145-153. ACM Press (2007)
2. Kempe, D., Dobra, A., Gehrke, J.: Gossip-based computation of aggregate information. In: Proc. of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 482–491. IEEE Press (2003)
3. Santoro, N., Sidney, J.B., Sidney S.J.: A distributed selection algorithm and its expected communication complexity. Theoretical Computer Science, vol. 100, pp. 185–204. Elsevier (1992)
4. Frederickson, G.N.: Tradeoffs for Selection in distributed networks. In: Proc. of the 2nd Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 154–160. ACM Press (1983)
5. Pardalos, P., Rajasekaran, S., Reif, J., Rolim, J.: Handbook on randomized computing. Kluwer Academic Publishers (2001)
6. Chaudhuri, S., Hagerup, T., Raman, R.: Approximate and exact deterministic parallel selection. LNCS, vol. 711, pp. 352–361. Springer (1993)
7. Ajtai, M., Komlós, J., Steiger, W.L., Szemerédi, E.: Optimal parallel selection has complexity $O(\log\log n)$. Journal of Computer and System Sciences, vol. 38, pp. 125–133. Academic Press (1989)

[4] $m^1_{d+1,j-1} + m^2_{d+1,j-1} \geq \Delta^{-2} \geq 1024$ so $m_{d,j} \geq \sum_{i=1}^{2} m^2_{d+1,j-1} > \frac{1024}{8} - \frac{1}{8} > 127$ by (2). Then $8\left(m_{d,j} + \frac{1}{8}\right) < 8\left(1 + \frac{1}{8\cdot 127}\right) m_{d,j} < 8.008$.

8. Frederickson, G.N., Johnson, D.B.: The complexity of selection and ranking in X + Y and matrices with sorted columns, J. Comput. System Sci., vol. 24 (2), pp. 197–208 Elsevier (1982).
9. Santoro, N., Suen, E.: Reduction techniques for selection in distributed files. IEEE Transactions on Computers, vol. 38, pp. 891–896. IEEE Press (1989)
10. Ajtai, M., Komlós, J., Szemerédi, E.: Sorting in $c \log n$ parallel steps. Combinatorica, vol. 3, pp. 1–19. Springer (1983)
11. Paterson, M.: Improved sorting networks with O(log N) depth. Algorithimica, vol. 5, pp. 65–92. Springer (1990)
12. Seiferas, J. L.: Sorting networks of logarithmic depth, further simplified. Algorithmica, vol. 53, pp. 374–384. Springer (2007)
13. Baddar, S. W. A., Batcher, K. E.: Designing sorting networks. Springer (2011)
14. Goodrich, M.T.: Zig-zag sort: a simple deterministic data-oblivious sorting algorithm running in O(n log n) time. In: Proc. of Symposium on Theory of Computing (STOC), pp. 684–693. ACM Press (2014)
15. Graham, R. L., Knuth, D. E., Patashnik, O.: Concrete mathematics. Addison-Wesley (1994)
16. Papadimitriou, C. H.: Computational Complexity. Addison-Wesley (1994)